

Sensor Fusion for Intelligent Behavior on Small Unmanned Ground Vehicles

G. Kogut*, G. Ahuja, B. Sights, E.B. Pacis, H.R. Everett
Space and Naval Warfare Systems Center, San Diego
53560 Hull Street, San Diego, CA 92152

ABSTRACT

Sensors commonly mounted on small unmanned ground vehicles (UGVs) include visible light and thermal cameras, scanning LIDAR, and ranging sonar. Sensor data from these sensors is vital to emerging autonomous robotic behaviors. However, sensor data from any given sensor can become noisy or erroneous under a range of conditions, reducing the reliability of autonomous operations. We seek to increase this reliability through data fusion. Data fusion includes characterizing the strengths and weaknesses of each sensor modality and combining their data in a way such that the result of the data fusion provides more accurate data than any single sensor. We describe data fusion efforts applied to two autonomous behaviors: leader-follower and human presence detection. The behaviors are implemented and tested in a variety of realistic conditions.

KEYWORDS: robotics, unmanned systems, data fusion, intelligent behaviors, computer vision

1. BACKGROUND

1.1 Technology Transfer Project

The JGRE Technology Transfer Project (TechTXFR) managed by Space and Naval Warfare Systems Center, San Diego (SSC San Diego) seeks to enhance the functionality (ability to perform more tasks) and autonomy (with less human intervention) of teleoperated systems¹. The objective is to expedite advancement of the technologies needed to produce an autonomous robot that can robustly perform in battlefield situations. Instead of developing new capabilities from scratch, the approach is to assess the technology readiness levels (TRLs) of component technologies (i.e., mapping, object recognition, motion-detection-on-the-move) developed under a variety of past and ongoing R&D efforts (such as the DARPA Tactical Mobile Robot program). The most mature algorithms are integrated and optimized into cohesive behavior architectures and then ported to various platforms used by the warfighter for further evaluation in operational environments.

Contributing sources of component technologies include the Idaho National Laboratory (INL), NASA's Jet Propulsion Laboratory, Carnegie-Mellon University (CMU), Stanford Research Institute International (SRI), University of Michigan, Brigham Young University, University of California San Diego, and University of Texas Austin, as well as other SSC San Diego projects (e.g., Man Portable Robotic System² and the ROBART series³). Starting in FY-03, the approach was to harvest existing indoor navigation technologies developed by various players and assess their different approaches to dead reckoning, obstacle detection/avoidance, mapping, localization, and path planning. The details of these focus areas will not be discussed in this paper but can be found in previous project publications⁴. The best features of the more promising solutions have now been integrated into an optimal system, giving an operator the ability to send an autonomous platform into an unknown indoor area and accurately map the surroundings. An *augmented virtuality* representation of the environment is derived, fusing real-time sensor information with the evolving map. In FY-05, the focus was expanded to include autonomous outdoor navigation, as well as additional sensor payloads for mission-specific applications such as intruder detection. As sensor technologies and autonomous behavior methods continue to be tested and evaluated in near-operational environments, the need for sensor fusion becomes readily apparent to provide a more robust solution to the warfighter.

*unmannedsystems@spawar.navy.mil; tel: 619-553-0707; fax: 619-553-6188; www.spawar.navy.mil/robots

Report Documentation Page			<i>Form Approved OMB No. 0704-0188</i>	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE 2007	2. REPORT TYPE	3. DATES COVERED 00-00-2007 to 00-00-2007		
4. TITLE AND SUBTITLE Sensor Fusion for Intelligence Behavior on Small Unmanned Ground Vehicles		5a. CONTRACT NUMBER		
		5b. GRANT NUMBER		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)		5d. PROJECT NUMBER		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Space and Naval Warfare Systems Center, San Diego, 53560 Hull Street, San Diego, CA, 92152		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)		
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 11
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified		

1.2 Intelligence Kernel

All component technologies described here are integrated under an expanded version⁴ of a robot architecture called the *Intelligence Kernel*, originally developed by INL⁶. To ensure cross-platform compatibility, the architecture is independent of the robot geometry and sensor suite, facilitating easy porting to any platform the warfighter uses. Moreover, the *Intelligence Kernel* allows the robot to recognize what sensors are available at any given time and adjust its behavior accordingly. The *Intelligence Kernel* facilitates the development of data fusion algorithms by abstracting and publishing all sensor data and derived data, called perceptions, in an easy-to-use manner.

2. HARDWARE DESCRIPTION

Two sensors were used in the experimentation described in this paper. Both sensors are off-the-shelf and are common to man-portable robots. However, because sensor set-up and use can vary greatly, a brief description on the hardware setup is provided.

2.1 Thermal Imager

The thermal camera used is the FLIR Systems *ThermoVision* A10. This imager was selected because its small size is appropriate for man-portable robots. The A10 is also well-suited for human presence detection in several other aspects. It uses a microbolometer detector, allowing it to measure absolute temperatures in a scene. This allows for easier segmentation of humans in the presence of objects hotter than the human. The A10 has a spectral response of 7.5 to 13.5 microns, which matches the peak wavelength light emitted by humans⁷.

The set up of a thermal camera for automated segmentation and detection is extremely important, though not always described in the literature of those who have performed automated detection and tracking with thermal imagery. Most off-the-shelf thermal cameras are designed so that the imagery output by their default settings closely resembles an image produced by a visible-light camera. This is not unexpected; most applications of thermal cameras involve direction interpretation of the imagery by humans who need the imagery in the most readily accessible form possible. Examples of such applications include thermal rifle-scopes and night-vision goggles for helicopter pilots. However, the signal processing to render thermal data so easily accessible to humans can greatly alter the raw data from the thermal detector and complicate subsequent image processing. The signal processing, called *Smart Scene* on the *ThermoVision* A10, introduces two unwanted effects. It tends to increase the intensity values of colder, background objects of the scene to make them more visible than they would otherwise be by scaling both cold and hot pixels in the scene so that they are both maximally visible in the resulting imagery. This has the effect of often reducing the contrast between hot and cold objects, making segmentation more difficult.

Another effect is that this contrast adjustment is dynamically adjusted in real-time depending on the thermal composition of the environment. The sudden introduction of an extremely hot object may result in the brightness values of the existing environment being suddenly scaled down so that the new object does not saturate the image. This effect can make segmentation difficult by forcing the image processing algorithm to alter thresholds in real-time to “keep up.”

These default signal processing steps are turned off in our application so that the thermal energy emitted by human skin results in approximately the same intensity value in the resulting imagery regardless of the surrounding environment. Examples of thermal imagery with and without *Smart Scene* are shown Figure 1. The image on the left “enhances” the visibility of the doorways and also shows some solar loading of a door in the lower left corner of the images with brightness levels that approach that of the human skin. The image on right shows the same scene with fixed intensity levels and no contrast enhancement. Notice that the image on the right is almost self-segmenting, greatly simplifying subsequent image processing steps and reducing possible false alarms.



Figure 1: Thermal Imagery from the *ThermoVision* A10 with and without onboard contrast enhancement processing.

2.2 LIDAR

The LIDAR used is a standard SICK LMS 220 LIDAR with a 180-degree field-of-view approximately 90m range, with \sim 10mm resolution and \sim 15mm systematic error. The SICK LIDAR is a standard component of many robots.

3. HUMAN PRESENCE DETECTION

Human presence detection is an important application in military and first-responder robotics. Many applications benefit from the ability to detect and locate human presence, including explosive-ordnance disposal, building exploration, and tactical applications. The TechTXFR project has prioritized human presence detection as one of the primary capabilities that could improve the capability of small robots. Prior work on human presence by SSC San Diego and INL demonstrated initial successes in limited environments but required the robot to construct a map of the surrounding environment before humans could be effectively detected⁸.

In indoor security applications, motion-detection is often used as a surrogate for human presence detection. However, in robotics, this is not feasible since the robots themselves are moving, making many motion-detection algorithms difficult to implement, and in most robotic applications, there is a potential for many non-human objects in the environment to be moving as well. Typical sensors used for human presence detection include Doppler radar and thermal cameras^{7,8}. However, these are subject to false alarms and multipart problems in some environments. New technologies, such as the microwave radiometer¹⁰ are promising but have not yet reached a maturity level suitable for deployment.

We focus on fusing data from a LIDAR, a thermal camera, and a color camera. We have found these sensors to be complementary in that they too have non-overlapping strengths and weakness, such that the combination of sensors makes a much stronger human presence device than any single sensor. These sensors are also useful because they commonly exist on many robots, allowing the addition of this human-detection system without the addition of expensive, specialized equipment. Table 1 below rates each sensor in a number of important measurements used in human presence detection. The last row in the table shows the capability of a theoretical “perfect” data fusion that could perfectly fuse the best aspects of each sensor. Of course we do not claim to have developed such a perfect algorithm, but present the table to highlight the potential of data fusion. Perfect fusion is difficult in this application because we are using heterogeneous sensors, imagers and laser, which provide fundamentally different data types. We describe a two-stage fusion process, called the *anomaly verification* process, which concentrates on using the strengths of each sensor.

	Human Signature Recognition	Depth Measurement	Range of Detection	Size Measurement	Motion Detection	Field-of-view
FLIR	Excellent	Weak	Variable	Excellent	Poor	Poor
LIDAR	Weak	Excellent	Excellent	Weak	Excellent	Excellent (2D)
Color Imagery	Mediocre	Weak	Mediocre	Excellent	Poor	Poor
Perfect Fusion	Excellent	Excellent	Excellent	Excellent	Excellent	Excellent

Table 1. Comparison of the strengths and weaknesses of sensors used.

3.1 Anomaly Detection

The first stage of human-presence detection is anomaly detection. An anomaly, in our application, is defined as one of two things: an entity not expected according to *a priori* map, or an object moving relative to the robot. In the latter case, the relative motion may also occur while the robot itself is moving. Anomaly detection is preformed by the LIDAR. The LIDAR is ideal for anomaly detection because of its range, accuracy, and large field of view. We use a common off-the-shelf 2D lidar, but 360-degree lasers and 3D flash LADARS will soon be available for testing. Anomalies are tracked and recorded, along with information about their position, size, and velocity. Two methods of anomaly detection are described. One method requires the construction of an occupancy map and the other requires no *a priori* information.

3.2 Anomaly Detection with a Map

This implementation relies on a simultaneous localization and mapping (SLAM) algorithm, developed by Kurt Konolige⁹ at Stanford Research Institute International (SRI) and optimized under the Technology Transfer Project, to characterize the environment and to maintain accurate localization of the robot as it navigates. SSC San Diego's partners under the Technology Transfer project, INL, leveraged SRI's LADAR based SLAM technology to develop a change-detection-on-the-move capability, called the INL Real Time Occupancy Change Analyzer (ROCA)¹⁰. This capability uses part of the occupancy grid from the SLAM algorithm to detect changes in the environment based on the robot's surrounding map grid. The changes are visible in Figure 3 as blue cubes in front of the robot. The location of the change is then sent as a vector to a supporting thermal imager for further assessment of human signature presence. In this implementation, though, the change detection only works if re-visiting an already mapped area. This is because the detection is done by finding differences between the current ladar scans and the known occupancy grid.

3.3 Anomaly Detection without a Map

This implementation takes advantage of a LADAR-based real-time environment feature extraction perception, data association, and tracking tools built into the Intelligence Kernel. The feature extraction capability uses LADAR scans to help find and define large changes in the environment as the robot drives through. The data association tool takes advantage of these changes to associate each new observation with the older ones. The data association method used is a variation on the K-nearest neighbor algorithm. Information about each change in the environment, e.g., location and size, are mapped into a multi-dimensional feature space. Incoming change detections are assigned to regions of this space relating to their proximity according to a distance measure and their class label. The distance measure used is the Mahalanobis distance. The Mahalanobis distance is preferable to Euclidean distance because features often vary greatly in measurement and systematic noise and should not be treated equally in a distance calculation. Other methods of data association and tracking, such as particle-filtering, will be explored in future work.

The tracked locations and location variances of moving objects, along the known position and velocity of the robot, are used to distinguish immobile objects from objects moving relative to the robot. Using the current location of the change in a world, as opposed to robot-centric coordinate system, we find the location in terms of the camera's orientation, which allows for moving the camera in that direction for verification of human presence. The detection is real-time,

requiring no prior knowledge of the area, and once the robot has explored and mapped an area, the robot automatically switches from this technique to ROCA and vice-versa. This method works well, but has a somewhat higher false alarm rate than the ROCA method. As of this writing, not enough data has yet been collected to characterize system performance.

3.4 Verification

Once an anomaly has been detected and localized, the next step is verifying human presence detection from a thermal camera. This is done by segmenting regions of temperature consistent with human presence from imagery. We followed the approach developed by Conaire which employs image histograms to select regions with temperatures likely to be produced by humans from the background¹¹. This method performs reliably in most environments. Sample segmentation images are shown in Figure 2. The figure on the left is a segmented image showing extractions of regions consistent with human skin temperature. The image on the right shows segment centroid calculation for the largest connected-component of the segmented image. This component is usually the head of a person, which usually has the largest area of exposed skin and, therefore, the largest thermal emission.



Figure 2: A segmented image (left) and centroid calculation (right).

Once the regions are located, their centroids and sizes are calculated, along with some shape descriptors, such as aspect ratio and degree of convexity. The shape descriptors are used to reject shape unlikely to have been produced by humans, such as perfect squares, etc. Finally, thresholds for temperature and size were used to eliminate noise and extraneous warm objects. These thresholds were calculated using a minimum squared-error optimization technique based on several hours of collected thermal imagery from a variety of indoor and outdoor scenes, temperatures, and weather conditions.

Figure 3 shows images of the leader-follower algorithm in action. The right image shows the robot following the leader outside from a building. The left image shows the perspective from the INL 3D Interface¹⁰.




Figure 2. The leader-follower behavior in action. In the left image, the blue block in front of the robot is generated by the presence of an obstacle. The thermal signature from the FLIR verifies human presence.

3.5 Results for the Verification Stage

The algorithm was tested with approximately 13,000 images from 2 hours of recorded imagery encompassing three scenarios: 1) indoor lab environment, 2) outdoor cold environment, 3) outdoor warm environment with significant solar loading of surfaces, and 4) indoor with non-human warm objects. The images were hand-classified into ground truth sets for images containing human presence and images without humans for purposes of generating detection metrics. All humans were within 25m of the thermal imager. The detection rates and false alarm rates are shown in Table 2.

	Indoor	Outdoor – cold	Outdoor – warm
Detection Rate	96%	92%	77%
False Alarm Rate	1%	3%	15%

Table 2: Detection and False Alarm Rates for human presence detection

As expected, warm exterior environments presented problems due to some surfaces being heated to temperatures comparable to human skin. However, these results are limited to the verification stage. These false alarm rates could be reduced considerably by only considering anomalies detected in the first stage of the fusion algorithm. Ad-hoc testing indicates lower false alarm rates in all environments for the fusion system. However testing is still underway and not reportable at the time of writing.

3.6 Color-Thermal Fusion

Still another form of fusion we are exploring for human-presence detection is more conventional fusion of color and thermal imagery. This technique has been used by Fujimasa¹¹ and others in medical imagery, and several others in human detection and tracking^{11, 12}. We employ similar fusion in verification of human presence. A common cue in color imagery used to detected humans is skin hue. The hue of human skin tone is relatively invariant to lighting conditions as well as the ethnicity of humans¹³. This invariance has made skin hue a useful tool in face detection algorithms¹⁴ and should also make a useful tool to aid in detecting human presence. However, because other objects with skin hue may exist in the environment, the false alarm rate of skin hue makes it too unreliable as a standalone presence sensor. However, if we register thermal imagery with color imagery, we can fuse their results and produce a detector that outperforms either of the individual detectors. Our fusion algorithm occurs at the pixel level and is based on the *general fusion model* described by Conaire¹¹ that calculates a probability of skin presence by a weighted combination of the fit to component models (skin hue model and thermal model). The weighting factor is important because it allows the verification system to prefer one model over the other depending on the range of the object, the application, and the environment. For example, thermal imagery may be unreliable in extremely hot, outdoor environments, while color imagery does not work well in very low light.

Sample images showing the image overlay technique are shown in Figure 4. The image on the left shows a thermal image overlaid directly on a color image. The imperfect mounting of our color camera results in the skewed overlay of the images. Regions which are likely to correspond to human skin or human thermal signature are highlighted in the fused image on the right. While initial results suggest that this method will both improve the detection rate and reduce the false-alarm rate of human presence detection, publishable results are not available at the time of writing.



Figure 3: Thermal-Color fusion

4. LEADER-FOLLOWER BEHAVIOR

A useful behavior for small robot operations is the leader-follower behavior. This is a mode of operation where a robot follows a person in a manner similar to a well-trained dog. This behavior prevents the need for operators to have to manually tele-operate or carry the robot while moving from place to place. Most implementations of this behavior employ GPS, as in the Jet Propulsion Laboratory's or SSC San Diego's leader-follower systems^{18, 19}. GPS works well outdoors but will not work indoors or in other GPS-denied areas. Vision-based methods can work well in some situations but do not provide accurate range (without stereo) and require clear visibility and good lighting to work properly. This application should also be distinguished from the large volume of research in large-vehicle convoying, which is a related, but fundamentally different task from small robot leader-follower. We demonstrate a system that fuses perceptions from three independent sensors: LIDAR, a thermal camera, and a monocular color camera. Each sensor tracks the leader independently and can be used alone with some success in the leader-follower behavior. However each also has specific weaknesses. Their output is intelligently fused so that when one sensor fails or provides noisy or weak data, the system will rely more heavily on the other sensors.

4.1 Algorithm Description

When following a target using the ladar, we mainly rely on two simultaneous algorithms. The first algorithm searches the field of view for the closest object. The size of the field of view and the center angle are adjusted depending on the previous distance from the robot to the target, the size of the target, and the predicted next position of the target. Adjusting the size keeps the algorithm from finding the edges of stationary objects as the target passes near them. The second algorithm uses laser edge perception to match the closest edges in the current laser scan to the predicted target location based on the previously calculated velocity vector of the target. The idea behind this algorithm is to keep the overall perception from being confused when a target follows along a wall or other object because even though the closest distance to the robot in the field of view may be the wall, the edges of the target still stand out (Figure 5).




Figure 5: The left image shows ladar based tracking looks for the minimum range and the edges of the target. The right image shows ladar-based and vision-based calculated angles to target when a person moves back and forth in front of the robot. Notice that the ladar data is much noisier due to the fact that it tracks whichever leg of the person is closest to the robot as they walk.

When following a target using the vision data, we rely on the pan angle of the camera and the location of the target in the camera image to calculate our heading error. The vision algorithm used is identical to that of section 2 of this paper, and detects, locates, and tracks the human leader. Calculating the range error with data from a monocular camera is slightly more difficult. The first method assumes our target is a certain height and, therefore, we can calculate our range based on the vertical offset of the camera on the robot, the tilt of the camera, and the vertical location of the target in the camera image. Alternatively, we can extrapolate a line along the camera's pan axis and then find the range sensor readings which are closest to that line. The first method has the advantage that the robot can know to go around obstacles if they get in the way instead of assuming that the closest range reading must be coming from the target.

4.2 Fusing the Outputs

The fusion method that we implemented involves the use of a fuzzy logic arbiter which takes the range and bearing to the target from the ladar and vision-based following algorithms, as well as the confidence in the measurement of each parameter (Figure 6). Calculating the confidence for each method is a critical part of the implementation as it helps the arbiter decide which algorithm to trust predominantly at any given moment. However, cases still arise whereby each algorithm believes it is correctly tracking the target but the algorithms disagree. In these cases, the Fuzzy Associative Memory (FAM) rules are designed to make up for weaknesses in each algorithm. For instance, one of the weaknesses of the ladar based tracking algorithm is that it occasionally decides that the edges of stationary objects are the desired target. In these cases, the robot will stop and face this object as the real target continues to move away. The FAM rules, in this case, say that if the LADAR-based perception is targeting an object that is "Very Close" and has a Zero angle to target, but the vision algorithms are targeting at a "Large" or "Very Large" angle to target, then the output yaw speed should reflect the desired yaw direction of the vision algorithms.




Figure 6: Diagram showing how the ladar and vision target range, angle, and confidence are fused using a fuzzy logic arbiter and then passed to the fuzzy logic obstacle-avoidance to determine the velocity vector that is sent to the drive control system.

4.3 Experimental Results

Overall, the ladar based method alone works very well and provides for a reasonably robust and aggressive behavior. However, when the target travels through small areas, such as doorways, or if the target follows the contours of walls or stationary objects, the ladar-based method can become confused. On the contrary, the vision-based methods work very robustly as the target passes through doorways and along walls, so fusing the data can produce a more robust behavior. The color-based vision tracking works decently, but does not transition well between indoor and outdoor environments. Finally, the FLIR-based vision tracking is robust and fusing it with the ladar data has created a very useful algorithm. The only drawback to the FLIR method is that it occasionally confuses the human target when there are several around.

6. CONCLUSION AND FUTURE WORK

As robotic behaviors become more complex, sensors will become increasingly important. Data fusion has been proven to increase the reliability of a system beyond that possible by multiple, independent sensors. While some data fusion tools, such as Kalman filters, have long been used, some applications and heterogeneous sensor types require unconventional methods of fusion. We've described the use of a fuzzy logic system and a 2-stage *anomaly-verification* method to increase the reliability of two useful behaviors for small robots. These methods require no special calibration steps or hardware not already commonly used on autonomous mobile robots.

Future work includes developing performance metrics for these behaviors and providing a detailed characterization of their performance in real-world environments. We also seek to develop a generalized motion-detection-on-the-move system for detecting and localizing moving objects while the robot itself is moving. Such an algorithm would be useful in many robot behaviors, such as pedestrian avoidance and target tracking.

7. REFERENCES

1. Pacis, E.B., Everett, H.R., Farrington, N., Kogut, G., Sights, B., Kramer, T., Thompson, M., Bruemmer, D., and D. Few, "Transitioning Unmanned Ground Vehicle Research Technologies," SPIE Proc. 5804: Unmanned Ground Vehicle Technology VII, Orlando, FL, March 29-31, 2005.
2. Bruch, M.H., Gilbreath, G.A., Muelhauser, J.W., and J.Q. Lum, "Accurate Waypoint Navigation Using Non-differential GPS," AUVSI Unmanned Systems 2002, Lake Buena Vista, FL, July 9-11, 2002
3. <http://www.spawar.navy.mil/robots/land/robart/robart.html>
4. Sights, B., Ahuja, G., Kogut, G., Pacis, E.B., Everett, H.R., Fellars, D., Hardjadinata, S., "Modular Robotic Intelligence System based on Fuzzy Reasoning and State Machine Sequencing," SPIE Proc. 6561: Unmanned Systems Technology VIII, Defense Security Symposium, Orlando, FL, April 97-13, 2007
5. Sights, B., Everett, H.R., Pacis, E.B., and G. Kogut, "Integrated Control Strategies Supporting Autonomous Functionalities in Mobile Robots," Computing, Communications, and Control Technologies Conference, Austin, TX, July 24-27, 2005.
6. D. J. Bruemmer, J. L. Marble, D. D. Dudenhoeffer, M. O. Anderson, and M. D. McKay. Mixed-Initiative Control for Remote Characterization of Hazardous Environments, *HICSS 2003*, Waikoloa Village, Hawaii, January 2003
7. Arlowe, H. D., "Thermal Detection Contrast of Human Targets," International Carnahan Conference, Atlanta, GA, October 14-16, 1992.
8. Kogut, G. T., Birchmore, F., Pacis, E. B., Ahuja, G., Sights, B., and B. Everett, "Using advanced computer vision algorithms on small mobile robots," SPIE Proc. 6230: Unmanned Systems Technology VIII, Defense Security Symposium, Orlando, FL, April 17-20, 2006.
9. Kurt Konolige, Ken Chou: Markov Localization using Correlation, *IJCAI 1999*: 1154-1159.
10. D.J. Bruemmer, D.A. Few, M.C. Walton, R.L. Boring, J.L. Marble, C. Nielsen, and J. Garner. "Turn off the television!: Real-world robotic exploration experiments with a virtual 3-D display." To appear in Proceedings of the Hawaii International Conference on System Sciences (HICSS) 2005.
11. C. Conaire, E. Cooke, N. O'Connor, N. Murphy, A. Smeaton, "Background Modeling in Infrared and Visible Spectrum Video for People Tracking", Proc. of IEEE Int'l workshop on Object Tracking & Classification Beyond the Visible Spectrum, 2005.
12. I. Fujimasa, A. Kouno, H. Nakazawa, Development of a new infrared imaging system: an infrared image superimposed on the visible image, 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Vol 20, No 2, 1998.
13. D. Borghys, P. Verlinde, C. Perneel, M. Achery, Multi-Level Data Fusion for the Detection of Targets using multi-spectral Image Sequences, SPIE Optical Engineering's special issue on Sensor Fusion, Vol 37, No 2, 477-484, 1998.
14. E. Saber and A. Tekalp. Frontal-View Face Detection and Facial Feature Extraction using Colour, Shape and Symmetry based Cost Functions. *Pattern Recognition Letters*, pages 669-680, 1998.
15. P. Lombardi, A Survey on Pedestrian Detection for Autonomous Driving Systems, Technical Report, University of Pavia, 2001.
16. C. Therrien, J. Scrofani, W. Krebs, An adaptive Technique for the Enhanced Fusion of Low-Light Visible with Uncooled Thermal Infrared Imagery, International Conference on Image Processing, 1997.
17. M. Jones and J. Rehg. Statistical Color Models with Application to Skin Detection. *CVPR*, page 274, 1999
18. Hogg, R., et al, "Sensors and algorithms for small robot leader/follower behavior," Proc. SPIE Vol. 4364, p. 72-85, Unmanned Ground Vehicle Technology III, Grant R. Gerhart; Chuck M. Shoemaker; Eds.
19. Nguyen, H.G., Kogut, G., Barua, R., Burmeister, A., Pezeshkian, N., Powell, D., Farrington, N., Wimmer, M., Cicchetto, B., Heng, C., and V. Ramirez, "A Segway RMP-based robotic transport system," SPIE Proc. 5609: Mobile Robots XVII, Philadelphia, PA, October 26-28, 2004

